

Nonnormal amplification in random balanced neuronal networks

Guillaume Hennequin

Tim P. Vogels

Wulfram Gerstner

School of Computer and Communication Sciences and Brain Mind Institute
Ecole Polytechnique Fédérale de Lausanne – 1015 EPFL, Switzerland

Correspondence: guillaume.hennequin@epfl.ch

Abstract

In dynamical models of cortical networks, the recurrent connectivity can amplify the input given to the network in two distinct ways. One is induced by the presence of near-critical eigenvalues in the connectivity matrix \mathbf{W} , producing large but slow activity fluctuations along the corresponding eigenvectors (dynamical slowing). The other relies on \mathbf{W} being nonnormal, which allows the network activity to make large but fast excursions along specific directions. Here we investigate the tradeoff between nonnormal amplification and dynamical slowing in the spontaneous activity of large random neuronal networks composed of excitatory and inhibitory neurons. We use a Schur decomposition of \mathbf{W} to separate the two amplification mechanisms. Assuming linear stochastic dynamics, we derive an exact expression for the expected amount of purely nonnormal amplification. We find that amplification is very limited if dynamical slowing must be kept weak. We conclude that, to achieve strong transient amplification with little slowing, the connectivity must be structured. We show that unidirectional connections between neurons of the same type together with reciprocal connections between neurons of different types, allow for amplification already in the fast dynamical regime. Finally, our results also shed light on the differences between balanced networks in which inhibition exactly cancels excitation, and those where inhibition dominates.

1 Introduction

A puzzling feature of cortical dynamics is the presence of structure in spontaneously generated activity states. For example, activity in cat primary visual cortex fluctuates along some non-random spatial patterns even when recordings are performed in complete darkness [1, 2]. Similarly, spontaneously generated patterns of firing rates in rat sensory cortices occupy only part of the total space of theoretically possible patterns [3]. As the constraints that govern these dynamics cannot be attributed to external stimuli, they are thought to originate from the patterns of synaptic connectivity within the network [4, 5]. This phenomenon is called patterned amplification.

Patterned amplification can also be observed in simulated neuronal networks, in which spontaneous activity can be modelled as the response to unspecific, noisy inputs delivered to each neuron individually. Propagated through recurrent connections, these noisy inputs may cause the activity of some neurons to transiently deviate from their average more strongly than could be expected from the variability of the external inputs. We thus define amplification here as the strength of these additional, connectivity-induced fluctuations.

Let us consider the following simple linear model for stochastic network dynamics:

$$d\mathbf{x} = \frac{dt}{\tau} (\mathbf{W} - \mathbb{1}) \mathbf{x} + \sigma_{\xi} d\boldsymbol{\xi} \quad (1)$$

where τ is the neuronal time constant, $\mathbf{x} \in \mathbb{R}^N$ is the deviation of momentary network activity with respect to a constant mean firing rate, \mathbf{W} is an $N \times N$ synaptic connectivity matrix, $\mathbb{1}$ is the identity matrix, and $d\boldsymbol{\xi}$ is a noise term modelled as a Wiener process. The fluctuations of $x_i(t)$ around zero (i.e. around the mean firing rate of neuron i) is caused by the noisy input and the recurrent drive. Starting from arbitrary initial conditions, the network activity \mathbf{x} converges to a stationary Gaussian process with covariance matrix $\boldsymbol{\Sigma} = \{\sigma_{ij}\}$ (at zero time lag), provided no eigenvalue of \mathbf{W} has a real part greater than unity. This covariance matrix has a baseline component $\boldsymbol{\Sigma}_{\text{unc.}} = \sigma_{\xi}^2 \tau \mathbb{1} / 2$ that corresponds to the covariance matrix in the absence of network connections (“unconnected”). Wiring up the network yields additional correlations and gives potentially rise to larger fluctuations of the activity of individual units. We define this amplification A as the ratio $[\text{Tr}(\boldsymbol{\Sigma}) - \text{Tr}(\boldsymbol{\Sigma}_{\text{unc.}})] / \text{Tr}(\boldsymbol{\Sigma}_{\text{unc.}})$. In other words, A measures the relative gain in mean variance that can be attributed to the recurrent connections. That is,

$$A(\mathbf{W}) \stackrel{\text{def}}{=} \left[\frac{2}{\tau \sigma_{\xi}^2 N} \sum_{i=1}^N \sigma_{ii} \right] - 1 \quad (2)$$

Under linear dynamics like that of Eq. 1, amplification can originate from two separate mechanisms. A first, “normal” type of amplification can arise from eigenvalues of \mathbf{W} with real parts close to (but smaller than) 1. The noise accumulates along the associated eigenvectors more than in

other directions, giving rise to larger activity fluctuations and substantial dynamical slowing along those axes. If the synaptic connectivity is normal in the mathematical sense ($\mathbf{W}\mathbf{W}^\dagger = \mathbf{W}^\dagger\mathbf{W}$), it is the *only* mechanism through which the network can amplify its input [5]. Indeed, if \mathbf{W} is normal, its eigenvectors form an orthonormal basis. The sum of variances in this eigenbasis is therefore equal to the sum of variances of the neuronal activities in the original equations. Since linear stability imposes that every eigenvalue of \mathbf{W} has a real part less than one, the activity along the eigenvectors can only decay following some initial perturbation. In other words, a stable *normal* linear system is contractive: no initial condition can transiently be amplified. If the matrix \mathbf{W} is not normal ($\mathbf{W}\mathbf{W}^\dagger \neq \mathbf{W}^\dagger\mathbf{W}$), another, *nonnormal* type of amplification can also contribute [5, 6, 7, 8]. The eigenvectors are no longer orthogonal to each other, and the apparent decay of the activity in the eigenbasis can hide a transient growth of activity in the neurons themselves. Such growth can only be transient, for stability requirements still demand that the activity decay asymptotically in time.

Purely nonnormal amplification that does not rely on dynamical slowing may be ideally suited for sensory cortices that need to track inputs varying on fast timescales (possibly of order τ). It has also been identified as a key mechanism for short-term memory of past inputs, for in certain circumstances, hidden feedforward dynamics enables the network to retain information about a transient stimulus for a duration of order $N\tau$ [7]. The presence of noise as in Eq. 1 could limit this memory duration to $\sqrt{N}\tau$ [6, 9], but this is still much longer than the time τ in which individual neurons forget their inputs.

The above considerations apply to purposely structured networks [6, 7, 5]. It is not clear, however, how much of this beneficial kind of amplification can be expected to arise in randomly connected networks of excitatory and inhibitory neurons, a ubiquitous model of cortical networks. Murphy and Miller [5] convincingly argued that nonnormal amplification should generally be a key player in the dynamics of balanced networks, i.e. when strong excitation interacts with equally strong inhibition and when neurons can be only excitatory or inhibitory but not of a mixed type. When the connectivity is dense, or at least locally dense, weak patterns of imbalance between excitation and inhibition can indeed be quickly converted into patterns in which neurons of both types strongly deviate from their mean firing rates. Here, we revisit nonnormal amplification in the context of *random* balanced networks. We derive an analytical expression for the purely nonnormal contribution to amplification in such networks. The analysis reveals a strong tradeoff between amplification and dynamical slowing, suggesting that the connectivity must be appropriately shaped for a network to simultaneously exhibit fast dynamics and patterned amplification.

2 Separating the effects of normal and nonnormal amplification

In the introduction, we have distinguished normal from non-normal amplification. The Schur decomposition (Fig. 1) – a tool from linear algebra – offers a direct way to assess the contributions of both mechanisms to the overall amount of amplification $A(\mathbf{W})$. Any matrix \mathbf{W} can be written as $\mathbf{U}^\dagger(\mathbf{\Lambda} + \mathbf{T})\mathbf{U}$ where $\mathbf{U} = \{u_{ij}\}$ is unitary, $\mathbf{\Lambda}$ is a diagonal matrix that contains the eigenvalues λ_k of \mathbf{W} , and $\mathbf{T} = \{t_{ij}\}$ is strictly lower-triangular¹ (Fig. 1a–c). The lines of \mathbf{U} are called the Schur vectors (or Schur modes) and are all orthogonal to each other. If this decomposition is to avoid complex numbers, $\mathbf{\Lambda}$ is only block-diagonal, with 2×2 blocks containing the real and imaginary parts of complex conjugate pairs of eigenvalues, and 1×1 blocks containing the real eigenvalues. Importantly, because the Schur basis \mathbf{U} is orthonormal, the sum of variances in the basis of the Schur vectors is equal to the sum of the single neuron activity variances. Thus, in order to compute $A(\mathbf{W})$, one can instead focus on the activity fluctuations in an abstract network whose units correspond to spatial patterns of neuronal activity (the Schur vectors) and interact with a connectivity matrix $\mathbf{\Lambda} + \mathbf{T}$ (Fig. 1a, right). This matrix is lower-triangular, so the abstract network is effectively feedforward. In the Schur network, unit i receives its input from all previous units $j < i$ according to the i^{th} row of \mathbf{T} . Without input, the activity of unit i decays at a speed governed by eigenvalue λ_i .

A network with a *normal* connectivity matrix would have only self-feedbacks ($\mathbf{T} = 0$), thus being equivalent to a set of disconnected units with a variety of individual effective time constants, reflecting dynamical slowing or acceleration. Amplification-by-slowing therefore arises from $\mathbf{\Lambda}$ (Fig. 1b), which summarises all the “loopiness” found in the original connectivity. Conversely, when $\mathbf{\Lambda} = 0$, all units share a common time constant τ (which is also the time constant of the actual neurons) and interact in a purely feedforward manner via matrix \mathbf{T} (Fig. 1c). We refer to this case as “purely nonnormal”, because the network is then free of the unique dynamical consequence of normality, namely a modification of the speed of the dynamics². “Purely nonnormal” amplification therefore arises from matrix \mathbf{T} that reveals the functional feedforward connectivity hidden in \mathbf{W} .

The latter situation ($\mathbf{\Lambda} = 0$) is the focus of this paper. By substituting \mathbf{W} with \mathbf{T} in Eq. 1 and subsequently calculating $A(\mathbf{T})$ as defined in Eq. 2, we intend to reveal the fraction of the total amplification $A(\mathbf{W})$ in the neuronal network that cannot be attributed to dynamical slowing, but only to transient growth. This constitutes a functional measure

¹Upper, not lower, -triangular \mathbf{T} is more common in the literature, but we prefer to keep the flow of information forward (from the 1st to the N^{th} Schur mode) for notational convenience in our calculations.

²Note that quantifying nonnormality can be done in a variety of ways, e.g. through several measures of “departure from normality” [8]. Our concept of “pure nonnormality” is therefore more specific to our particular purpose, in that it expresses the absence of normal effects on the dynamics of the neurons.

of nonnormality. We carry out this analysis in a statistical sense, by calculating the expected amount of purely nonnormal amplification $\langle A(\mathbf{T}) \rangle$ where the average $\langle \cdot \rangle$ is over the random matrix \mathbf{W} . In section 3, the ensemble statistics of \mathbf{W} are defined, and related to the statistics of the non-zero entries of \mathbf{T} . In sections 4 and 5, we perform the calculation of $\langle A(\mathbf{T}) \rangle$.

3 Schur representation of neural connectivity matrices

Prior to calculating the nonnormal contribution to amplification in realistic neural connectivity matrices, we first analyse the statistical properties of the Schur triangle \mathbf{T} derived from a neuronal network where every pair of neurons has a certain probability of being connected in either direction. Specifically, we consider networks of $N/2$ excitatory and $N/2$ inhibitory neurons, with connectivity matrices \mathbf{W} drawn as follows³ (Fig. 1a):

$$w_{ij} = \frac{1}{\sqrt{N}} \cdot \begin{cases} +w_0 & \text{if } j \leq N/2 \\ -w_0 & \text{if } j > N/2 \\ 0 & \text{with proba. } (1-p) \end{cases} \quad \text{with proba. } p \quad (3)$$

Excitation and inhibition are thus globally balanced. The $1/\sqrt{N}$ scaling ensures that in the limit of large N , the eigenvalues $\{\lambda_k\}$ of \mathbf{W} become uniformly distributed inside the disk of radius

$$R = w_0 \sqrt{p(1-p)} \quad (4)$$

and centered around zero in the complex plane (Fig. 1b), with the exception of a few outliers [10]. To push the outliers inside the disk, we enforce that excitatory and inhibitory synapses cancel each other precisely for each receiving neuron, i.e. $\mathbf{W}\mathbf{v} = 0$ with $\mathbf{v} = (1, 1, \dots, 1)/\sqrt{N}$ [10, 11]. This constraint is also essential to the identification of the ensemble statistics of \mathbf{T} as detailed below. Such a “global balance” can be achieved by a Hebbian form of synaptic plasticity at inhibitory synapses in random spiking networks [12]. Here we enforce it by subtracting the row average (a small number) from every row (which accounts for the barely visible horizontal stripes in \mathbf{W} of Fig. 1a).

The main point in relating the statistics of \mathbf{T} to that of \mathbf{W} is to note that the Schur basis is unitary, so that the sum of squares in \mathbf{W} is also equal to the sum of squares in $\mathbf{\Lambda} + \mathbf{T}$. Thus

$$\sum_{1 \leq i, j \leq N} w_{ij}^2 = \sum_{1 \leq k \leq N} |\lambda_k|^2 + \sum_{i > j} t_{ij}^2 \quad (5)$$

From our choice of the weights w_{ij} (Eq. 3) and assuming that N is large enough, we can derive $\sum w_{ij}^2 \simeq Npw_0^2$. Furthermore, knowing that the eigenvalues lie uniformly inside the disk of radius R , we can write $\sum |\lambda_k|^2 \simeq NR^2/2$ which is

³It is straightforward to allow for any distribution of non-zero weights; as it turns out, this Dirac delta distribution achieves maximum nonnormal amplification.

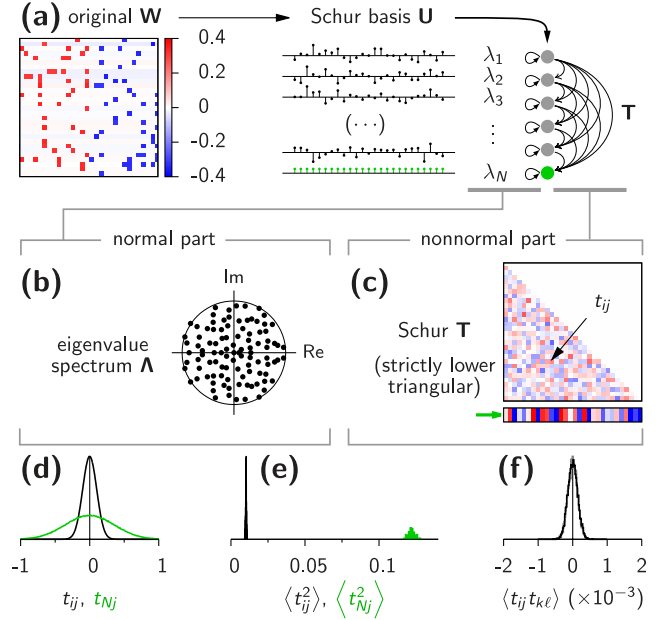


Figure 1: Teasing apart normal and nonnormal amplification in random networks of excitatory and inhibitory neurons. **(a)** Example sparse neural connectivity matrix \mathbf{W} (left, 50 exc. columns and 50 inh. columns, thinned out to 30×30 for better visibility), a schematics of an associated Schur basis \mathbf{U} (center), and the corresponding abstract network of Schur modes, in which the interactions are feedforward from top to bottom (right). The Schur vectors in \mathbf{U} (center), orthogonal to one another, represent patterns of neuronal activity in the original network. The last Schur vector is explicitly chosen to be the uniform “DC” mode $\mathbf{v} = (1, 1, \dots, 1)/\sqrt{N}$ and is represented here in green. **(b)** Amplification via dynamical slowing (“normal” amplification) is described by the set of eigenvalues $\mathbf{\Lambda} = (\lambda_1, \dots, \lambda_N)$ of \mathbf{W} , which for a random network lies inside a disk centered around zero in the complex plane. These eigenvalues determine the decay rates of the Schur patterns. **(c)** Nonnormal amplification arises from the strictly lower-triangular matrix \mathbf{T} which describe the purely feedforward part of the interactions between the Schur patterns. The first non-zero entry in the upper left corner of \mathbf{T} is t_{21} and represents the “forward” coupling from the first Schur mode onto the second. The last row $(t_{N1}, t_{N2}, \dots, t_{N(N-1)})$, zoomed-in (green arrow), is the coupling from the first $N-1$ Schur modes onto the last (uniform) Schur mode \mathbf{v} . **(d)** For a fixed large matrix \mathbf{W} , the non-zero entries t_{ij} in matrix \mathbf{T} are approximately normally distributed with zero mean and variance given by Eq. 9 (black histogram, for $j < i < N$). The entries in the last row have larger variance given by Eq. 8 ($i = N$, green histogram). **(e)** Moreover, the variance $\langle t_{ij}^2 \rangle$ across many realisations of \mathbf{W} is the same for all $j < i < N$ (black histogram). Similarly, $\langle t_{Nj}^2 \rangle$ is the same for all $j < N$ (green histogram). **(f)** The correlations $\langle t_{ij} t_{k\ell} \rangle$ (for $i \neq k$ or $j \neq \ell$) are negligible, as seen from a comparison of their empirical distribution (black) with surrogate data from triangular matrices in which non-zero entries are all drawn i.i.d. (grey, barely visible under the black curve). The data for panels (d-f) was acquired by Schur-transforming 5,000 random weight matrices of size $N = 100$, drawn as described in section 3 with connection density $p = 0.1$ and spectral radius $R = 1$.

also valid for large N . We replace these sums in Eq. 5, sim-

ply the result using Eq. 4, and obtain the overall empirical variance of the non-zero entries in \mathbf{T} , to leading order in N :

$$\frac{2}{N(N-1)} \sum_{i>j} t_{ij}^2 \simeq \frac{R^2}{N} \cdot \frac{1+p}{1-p} \quad (6)$$

Note that this empirical variance is not necessarily equal to the *ensemble* variance $\langle t_{ij}^2 \rangle - \langle t_{ij} \rangle^2$ for fixed i and j . In fact, we have observed that if the non-unique Schur basis is chosen arbitrarily, $\langle t_{ij}^2 \rangle$ computed over many realisations of \mathbf{W} is not uniform across rows, but rather tends to increase with row index i . This heterogeneity is difficult to characterise, and undermines the calculation of amplification developed in the next section. Fortunately, we can circumvent this problem by choosing the uniform eigenvector \mathbf{v} of \mathbf{W} as the last Schur vector: $u_{Nk} = 1/\sqrt{N}$ for all k ⁴. Coefficient t_{ij} then becomes distributed with the same zero mean and variance ζ^2 for all $j < i < N$, with the exception of the t_{Nj} coefficients which have higher variance ζ_0^2 (black and green lines in Fig. 1d and 1e, empirical observation). Note also that the ensemble pairwise correlations between coupling strengths t_{ij} and $t_{k\ell}$ with $i \neq j$ or $j \neq \ell$ seem negligible (Fig. 1f).

We now proceed in two steps. First, we focus on the variance of the elements in the *last* row of the Schur matrix \mathbf{T} , and then we turn to all the other non-zero components. To calculate variance $\zeta_0^2 = \langle t_{Nj}^2 \rangle$ we use the definition of \mathbf{T} and write for $j < N$

$$\begin{aligned} t_{Nj} &= \sum_{\ell=1}^N \sum_{k=1}^N u_{Nk} w_{k\ell} u_{j\ell} \\ &= \frac{1}{\sqrt{N}} \sum_{\ell=1}^N \left(\sum_{k=1}^N w_{k\ell} \right) u_{j\ell} \end{aligned} \quad (7)$$

To leading order in N we can write $\sum_k w_{k\ell} = \pm p w_0 \sqrt{N}$ where the \pm sign depends on ℓ being smaller than $N/2$ (+, excitatory) or greater (−, inhibitory) – see Fig. 1a. For $j < N$, the j^{th} Schur vector \mathbf{U}_j is orthogonal to the last Schur vector $\mathbf{v} \propto (1, 1, \dots, 1)$, so its components strictly sum to zero: $\sum_{\ell} u_{j\ell} = 0$. Moreover, because of the normalization, $\sum_{\ell} u_{j\ell}^2 = 1$. We can therefore approximate $u_{j\ell}$ by a stochastic process with zero mean and variance $1/N$. Assuming the $u_{j\ell}$ are uncorrelated, the variance of t_{Nj} is thus simply $w_0^2 p^2$ to leading order, which according to Eq. 4 is also

$$\langle t_{Nj}^2 \rangle \equiv \zeta_0^2 = \frac{R^2 p}{1-p} \quad (8)$$

Notably, the variance ζ_0^2 in the last row of coupling matrix \mathbf{T} is of order 1, and depends super-linearly on the connectivity density p (Fig. 2, green lines).

We now turn to the other rows $i < N$ of the Schur matrix \mathbf{T} . Because all components t_{ij} for $j < i < N$

⁴This is always possible, since a Schur basis can be constructed through Gram-Schmidt orthonormalisation of the eigenbasis of \mathbf{W} , so choosing \mathbf{v} to enter the process first results in \mathbf{v} being the last vector in a basis that makes \mathbf{W} lower-triangular

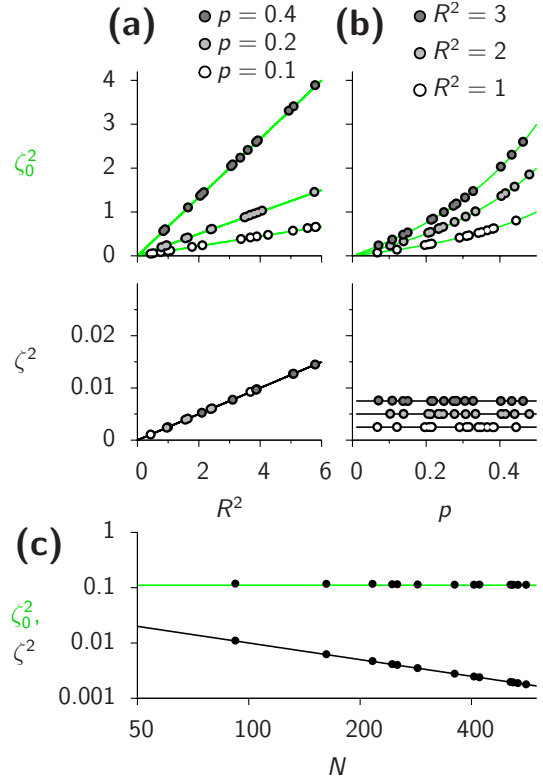


Figure 2: Linking the Schur triangle to the parameters of the neural connectivity matrix. (a) The variance of the entries in the strict lower triangle \mathbf{T} scales linearly with the square of the spectral radius R^2 of the original weight matrix \mathbf{W} . For the last row of \mathbf{T} , the slope of ζ_0^2 depends on the connection probability p (top plot). For the rest of \mathbf{T} , ζ^2 depends only on R^2 (bottom plot). Each point was obtained by empirically estimating ζ^2 and ζ_0^2 from 10 different Schur-transformed random neural weight matrices of size $N = 400$. Lines denote the analytical expressions in Eqs. 8 and 9. (b) ζ_0^2 in the last row of \mathbf{T} scales super-linearly with the connection density p (top plot). In contrast, ζ^2 does not depend on p (bottom plot). (c) In the last row of \mathbf{T} , the variance is network size-independent (green line). In the rest of \mathbf{T} , the variance is inversely proportional to N (black line, note the log-log scale).

seem to come from the same distribution and look uncorrelated (Fig. 1d–f), the empirical estimate of their variance $2 \sum_{j<i<N} t_{ij}^2 / ((N-1)(N-2))$ coincides with the ensemble variance $\zeta^2 \equiv \langle t_{Nj}^2 \rangle$ so long as N is large enough. Similarly, we can write $\sum_j t_{Nj}^2 / (N-1) = \zeta_0^2$. Thus, the l.h.s. of Eq. 6 becomes $\zeta^2 + 2\zeta_0^2/N$ to leading order in N . Using Eqs. 6 and 8 we conclude

$$\langle t_{ij}^2 \rangle \equiv \zeta^2 = \frac{R^2}{N} \quad (9)$$

Figure 2 shows that Eqs. 8 and 9 provide a good match to numerical results.

At this point we can already draw a few conclusions. Suppose each unit in our Schur network receives external input of variance 1. First, since the uniform mode \mathbf{v} receives network input from the remaining $N - 1$ Schur patterns with

coupling coefficients of order 1 (Eq. 8), we expect the global (“DC”) population activity $\mathbf{x} \cdot \mathbf{v}$ to fluctuate macroscopically, i.e. with a variance of order N . In contrast, the rest of the Schur modes should display fluctuations of order 1. Second, we directly see that making the network denser (i.e. increasing p) can only result in larger DC fluctuations, but no further amplification of the other (zero-mean) Schur patterns. This is because ζ_0^2 , but not ζ^2 , depends on p . Third, it is easy to see where these large DC fluctuations would originate from. Imagine breaking the overall exc.-inh. balance in the network activity by a small amount, e.g. by initialising the network state \mathbf{x} to $\mathbf{d} = (1, \dots, 1, -1, \dots, -1)/\sqrt{N}$, where we emphasize the scaling in $1/\sqrt{N}$. According to Eq. 1, the transient response to this perturbation is roughly $\mathbf{W}\mathbf{d}$, which to leading order in N equals

$$\mathbf{W}\mathbf{d} \simeq pw_0(1, 1, \dots, 1) \quad (10)$$

We note that the $1/\sqrt{N}$ scaling is gone. Thus, the network responds to a microscopic global balance disruption – a state in which the deviation between the excitatory and inhibitory population firing rates is of order $1/\sqrt{N}$ – by an excursion of order 1 in the combined firing rate of both populations (see [5] for a more in-depth discussion of this effect). Finally, it is instructive to see what happens when the functional feedforward link from \mathbf{d} to $\sqrt{N} \cdot \mathbf{v}$ – expressed in Eq. 10 – is removed from \mathbf{W} . This can be achieved by transforming \mathbf{W} into \mathbf{W}' given by

$$\mathbf{W}' = \mathbf{W} - \frac{pw_0}{\sqrt{N}}(1, \dots, 1)^\dagger(1, \dots, 1, -1, \dots, -1) \quad (11)$$

It is easy to see that $\mathbf{W}'\mathbf{d} = 0$. In this case, calculations similar to Eqs. 5–8 yield $\zeta_0^2 = \zeta^2 = R^2/N$ so that the DC fluctuations are back to order 1: the amplification along the DC mode becomes comparable in magnitude to the amplification that occurs along any other Schur directions. Note that the operation in Eq. 11 effectively shifts the mean excitatory (resp. inhibitory) weight from pw_0/\sqrt{N} (resp. $-pw_0/\sqrt{N}$) to zero. We now substantiate these preliminary conclusions through a direct calculation of nonnormal amplification.

4 Amplification in random strictly triangular networks

We have seen in the preceding two sections that a randomly coupled network of excitatory and inhibitory neurons can be transformed via a unitary Schur basis into a different network where the couplings between units are given by a lower triangular matrix (Fig. 1a). Furthermore, the “purely nonnormal” part of the amplification of the external noisy input in the original network of neurons corresponds to the activity fluctuations in the new feedforward network where all self-couplings are neglected (Fig. 1c). Finally, we have also seen that it is possible to constrain the Schur basis such that the couplings between the first $N - 1$ units in the feedforward network are independently distributed with the

same zero-mean and a variance given by the parameters of the original synaptic weights (Eq. 9). In this section, we therefore study this “canonical” case, starting directly from a strictly lower-triangular matrix \mathbf{T} and ignoring – for the moment – the transformation that gave rise to \mathbf{T} .

We want to solve for the expected variances of $N \gg 1$ Ornstein-Uhlenbeck processes (as in Eq. 1) coupled by a strictly lower-triangular weight matrix \mathbf{T} (therefore describing a purely feedforward network, see inset in Fig. 3a). We assume all non-zero coupling strengths to be sampled i.i.d. from some common distribution with zero mean and variance α^2/N . Due to the coupling matrix, the fluctuations that the external input causes in the first unit feed and augment those it causes in unit 2. The third unit in turn fluctuates due to the external input and the activities of units 1 and 2, and so on. We therefore expect the activity variance σ_{ii} in unit i to increase with index i . In appendix A, we show that in the limit of large N and for some fixed $0 \leq x \leq 1$, the relative expected variance of the activity in unit $i = xN$ is $g(i/N) \equiv 2 \langle \sigma_{ii} \rangle / \tau \sigma_\xi^2$ where the function $g(x)$ is lower-bounded in closed form by

$$g^{\text{LB}}(x) = \frac{1}{3 + \sqrt{3}} \exp\left(\frac{1 - \sqrt{3}}{4} \alpha^2 x\right) + \frac{2 + \sqrt{3}}{3 + \sqrt{3}} \exp\left(\frac{1 + \sqrt{3}}{4} \alpha^2 x\right) \quad (12)$$

(Fig. 3, dashed blue curves). We also derive the exact solution as a power series

$$g(x) = \lim_{K \rightarrow \infty} \sum_{k=0}^K \beta_k x^k \quad (13)$$

with the β_k coefficients defined recursively as

$$\beta_0 = 1, \quad \beta_k = \frac{\alpha^2}{2k!} \sum_{\ell=0}^{k-1} \frac{(2\ell)! (k - \ell - 1)!}{\ell! (\ell + 1)!} \left(\frac{\alpha^2}{4}\right)^\ell \beta_{k-\ell-1} \quad (14)$$

The overall amplification $A_0(\alpha^2)$ in the network is subsequently obtained by integrating this variance profile $g(x)$ from 0 to 1, which corresponds to taking Eq. 2 to its $N \rightarrow \infty$ limit:

$$A_0(\alpha^2) = \left(\lim_{K \rightarrow \infty} \sum_{k=0}^K \frac{\beta_k}{k+1} \right) - 1 \quad (15)$$

Figure 3 shows that Eqs. 13 and 15 indeed converge to the empirical mean variance profile and mean amplification as the cut-off parameter K of the power series becomes large (red lines, $K = 10$). Figure 3b furthermore shows how amplification explodes with the variance α^2/N of the feedforward couplings in the network.

5 Amplification in random balanced networks

Using the canonical result of the previous section that is restricted to homogeneous random lower-triangular matrices




Figure 3: Analytical result for a feedforward network of N Ornstein-Uhlenbeck processes coupled via a random strictly lower-triangular matrix (inset). **(a)** The expected activity variance $\langle \sigma_{ii} \rangle$ accumulates super-linearly from the first unit to the last down the feedforward chain. Dashed blue lines depict the closed-form lower-bound of Eq. 12. Red lines denote the exact solution given in Eq. 13, truncated to $K = 10$. Open circles represent the numerical solution of Eq. 1 – or more exactly, the numerical solution of equation 20 given in the appendix – averaged over 20 randomly generated matrices of size $N = 500$. Each matrix \mathbf{T} is characterised by the variance α^2/N of the coupling coefficients t_{ij} with $j < i$. The strength of the external noise driving each unit independently is set to $\sigma_\xi^2 = 2/\tau$ so that all activity variances in the network would be 1 should the couplings t_{ij} be set to 0. **(b)** The total amplification (the area under the curves in (a), minus 1) explodes with increasing variance α^2/N in the triangular connectivity matrix. Points and lines have the same meaning as in (a).

ces, we now calculate $A(R, p) \equiv \langle A(\mathbf{T}) \rangle$ with \mathbf{T} originating from the Schur decomposition of a neuronal connectivity matrix as in section 3, with connection density p and spectral radius R . Equation 13 can directly be applied with $\alpha^2/N = \zeta^2 = R^2/N$ (see Eq. 9) to describe the activity fluctuations of the first $N - 1$ Schur modes. The last Schur unit, however, receives feedforward input with couplings of variance $\zeta_0^2 \neq \zeta^2$ (Eq. 8). Consequently, the expected variance $\langle \sigma_{NN} \rangle$ of its temporal fluctuations has to be treated separately. In appendix B, we show that

$$\lim_{N \rightarrow \infty} \frac{\langle \sigma_{NN} \rangle}{N} = \frac{\sigma_\xi^2 \tau}{2} \cdot \frac{p}{1-p} [g(1) - 1] \quad (16)$$

where g is given by Eqs. 13 and 14, here with $\alpha = R$. Gathering the contributions of all Schur modes, we obtain the expected overall amount of purely nonnormal amplification

in \mathbf{W} :

$$A(R, p) = A_0(R^2) + \frac{p}{1-p} [g(1) - 1] \quad (17)$$

with $A_0(R^2)$ given by Eq. 15.

Figure 4a shows that the nonnormal contribution to amplification in the neuronal network explodes with the spectral radius R of the connectivity matrix \mathbf{W} . This is because the amplification of the first $N - 1$ Schur units explodes with the variance ζ^2 of their feedforward interactions (Fig. 3b) and that ζ^2 is directly related to R (Eq. 9). Note that for $R > 1$ (to the right of the dashed vertical line in Fig. 4a), the network of neurons is unstable. Although the concept of amplification in an unstable network is ill-defined, the “purely nonnormal” part of the total (infinite) amplification remains bounded. Indeed, the purely feedforward network \mathbf{T} derived from the Schur decomposition of \mathbf{W} is itself always stable, since zero is the only eigenvalue of \mathbf{T} . The instability in \mathbf{W} arises from purely normal effects, when the real part of one eigenvalue of \mathbf{W} exceeds unity so that dynamical slowing becomes infinite.

Equation 16 confirms what we had previously discussed at the end of section 3: the last Schur unit has temporal fluctuations $\mathbf{v} \cdot \mathbf{x}(t)$ of variance $\mathcal{O}(N)$. Those fluctuations thus make up for a finite fraction of the total nonnormal amplification (the last term in Eq. 17) as $N \rightarrow \infty$. Because the last Schur vector is the normalised uniform spatial pattern $(1, \dots, 1)/\sqrt{N}$, the variance of the overall population activity $\mu(t) \equiv \sum x_i(t)/N = \sqrt{N}(\mathbf{x} \cdot \mathbf{v}(t))$ is of order 1. As we had foreseen in section 3, one can restore the $1/N$ scaling of these “DC” fluctuations $\langle \mu^2(t) \rangle$ by performing the operation of Eq. 11 on the connectivity matrix \mathbf{W} , i.e. subtracting a common constant from all excitatory weights (including zero weights) to make sure that they average to zero, and adding the same constant to all inhibitory weights with the same purpose. This situation is depicted by the grey curves in Fig. 4. Figure 4b shows that only these DC fluctuations depend on the connectivity density p .

Overall, Fig. 4a allows us to draw two important conclusions. On the one hand, if the level of dynamical slowing is to be kept low ($R \ll 1$), only modest levels of amplification can be achieved (see the small amount of nonnormal amplification on the l.h.s. of the dashed vertical line). For example, if no mode is to decay with more than twice the single neuron time constant ($\text{Re}(\lambda) < 1/2$), the average variance cannot exceed that of a disconnected network by more than 10%. On the other hand, the nonnormal contribution to amplification explodes with increasing R , i.e. with increasing synaptic strengths if the connection density is taken fixed. This suggests that strong transient amplification without dynamical slowing can only be achieved in structured, “less random” networks. The structure must allow the synaptic couplings to assume larger values without causing the eigenvalue spectrum of \mathbf{W} to reach instability.




Figure 4: Nonnormal amplification in random neuronal networks. **(a)** The mean amount of purely nonnormal amplification $\langle A(\mathbf{T}) \rangle \equiv A(R, p)$ is reported as a function of the spectral radius R of \mathbf{W} . Open circles denote the numerical solution of Eq. 20 averaged over 20 randomly drawn connectivity matrices with connection density $p = 0.1$ and size $N = 500$. Errorbars denote the standard deviation over all trials. Red lines depict the exact solution in Eq. 17. The grey line and grey circles indicate the mean removal of Eq. 11 applied to \mathbf{W} , which effectively removes the global macroscopic fluctuations of the entire population (labelled “no DC”). The dashed vertical line represents the limit of linear stability, beyond which the nonnormal part of amplification is still well-defined. **(b)** Same as in (a), now as a function of the connection density p for a fixed $R = 1$. In both (a) and (b), parameters p and R fully determined the value $\pm w_0/\sqrt{N}$ of the nonzero synaptic weights as $w_0 = R/\sqrt{p(1-p)}$ (cf. Eq. 4).

6 Different numbers of excitatory and inhibitory neurons

We now consider the biologically more plausible case of different numbers of excitatory and inhibitory neurons. Typical models of cortex assume fN excitatory neurons and $(1-f)N$ inhibitory neurons with $f = 0.8$ or similar. In this case, the eigenvalues λ are no longer uniformly scattered inside the disk of radius R in the complex plane⁵, but become more concentrated in the middle following a radially symmetric density $\rho(|\lambda|)$ known analytically from [10] (Fig. 5b, insets). As before, we consider the case where excitatory (resp. inhibitory) synaptic couplings are 0 with probability

⁵Rajan and Abbott showed that this happens when the variances of the excitatory and inhibitory weights differ (the variances comprise both the zero and non-zero synapses). Decreasing the number of inhibitory neurons in a balanced network requires the strength of inhibition to be increased. In sparse networks like ours, this automatically makes the overall variance of the inhibitory synapses larger than that of excitatory synapses, hence the observed effect on the eigenspectrum.

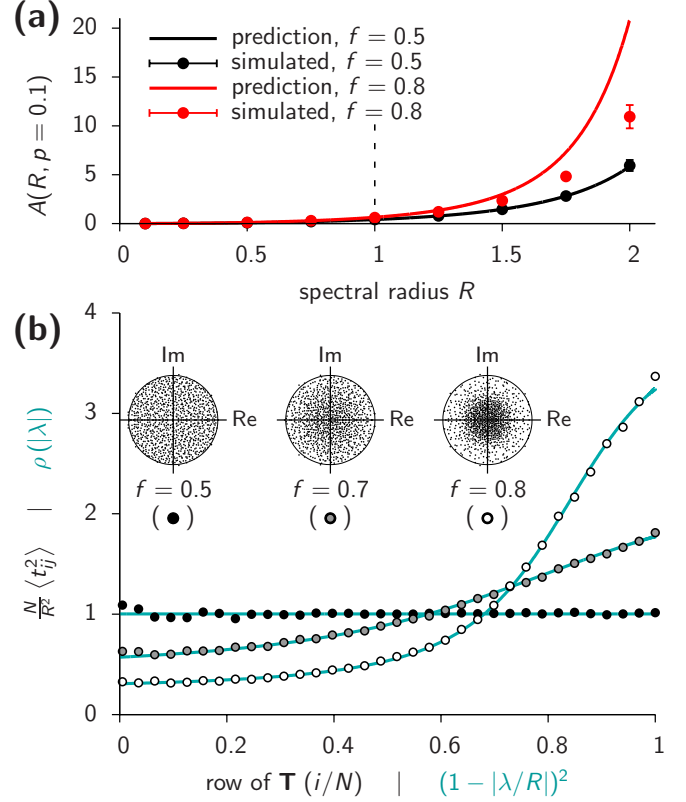


Figure 5: Networks with different numbers of excitatory and inhibitory neurons. **(a)** Nonnormal amplification as a function of the spectral radius R of \mathbf{W} , in sparse random balanced networks with fN excitatory and $(1-f)N$ inhibitory neurons, for $f = 0.5$ (black) and $f = 0.8$ (red). The connection density p was set to 0.1. The dashed vertical line represents the limit of linear stability, beyond which the nonnormal part of amplification is still well-defined. Solid circles were obtained by averaging the numerical solution of Eq. 20 for 20 random matrices of size $N = 500$. Errorbars denote standard deviation over all trials. **(b)** Filled circles show the scaled variance $N \langle t_{ij}^2 \rangle / R^2$ of the non-zero Schur couplings in row i as a function of i/N and for three different values of f . These variances were computed by Schur-transforming 100 matrices of size $N = 200$, with $R = 1$ and $p = 0.1$. Cyan lines denote the density ρ of eigenvalues λ inside the unit disk [10], as a function of $(1 - |\lambda/R|)^2$. Insets show the eigenvalue spectra of three example matrices of size $N = 1000$.

$(1-p)$, and $+w_E/\sqrt{N}$ (resp. $-w_I/\sqrt{N}$) otherwise. The global balance condition reads $fw_E = (1-f)w_I$. To impose a given spectral radius R , we set $w_E^2 = w_0^2(1-f)/f$ and $w_I^2 = w_0^2 f/(1-f)$ with $w_0^2 = R^2/p(1-p)$.

The results of section 3 regarding the variances in the Schur triangle have to be adjusted to accommodate these modifications. The derivation of ζ_0^2 is left unchanged, so that the couplings t_{Nj} onto the uniform mode \mathbf{v} still have the variance given by Eq. 8, which notably does not depend on f . Using Eq. 5, we can then write down the empirical variance in the first $N-1$ rows of \mathbf{T} as

$$\frac{2}{N(N-1)} \sum_{j < i < N} t_{ij}^2 = \frac{2}{N} \left(R^2 - \int_0^R r \rho(r) dr \right) \quad (18)$$

Unfortunately, the *ensemble* variance $\langle t_{ij}^2 \rangle$ for fixed i and j is in general different from the average across matrix elements given by Eq. (18). Indeed, contrary to the case $f = 0.5$ considered in section 3, the non-zero elements of \mathbf{T} *no longer* have the same ensemble variance. Instead, $\langle t_{ij}^2 \rangle$ grows with row index $i < N$, and this profile interestingly matches the density of eigenvalues ρ ⁶, according to

$$\frac{N}{R^2} \langle t_{ij}^2 \rangle = \rho \left[R \left(1 - \sqrt{\frac{i}{N}} \right) \right] \quad \text{for } j < i < N \quad (19)$$

This is depicted in Fig. 5b.

In a feedforward network like that of Schur units considered here, a good strategy to generate greater amplification would be to give comparatively more power to the couplings onto earlier nodes. This is because amplification builds up superlinearly along the feedforward chain (Fig. 3), so that boosting early nodes exacerbates the avalanche effect (see also [6]). Setting f to more than 0.5 does precisely the contrary: couplings onto early nodes become comparatively smaller in magnitude, as shown by the filled circles in Fig. 5b. Therefore, simply replacing α^2/N in Eq. 15 by the empirical variance of Eq. 18 yields an over-estimation of the true amplification in the first $N-1$ Schur units (compare the red line with the red circles in Fig. 5a). We found it difficult to incorporate this variance profile $\langle t_{ij}^2 \rangle$ into the derivation of appendix A, so we can only consider as accurate the results of numerical simulations.

The conclusions reached at the end of section 5 do not change significantly under the more realistic assumption of $f = 0.8$. Although amplification almost doubles relative to $f = 0.5$, it remains very weak in the stable regime (to the left of the dashed vertical line in Fig. 5a), confirming that amplification can only come with substantial dynamical slowing when connections are drawn at random.

7 Example of network structure for nonnormal amplification

Here we show that random networks can be minimally structured in such a way that strong nonnormal amplification occurs already in the fast dynamical regime. We exploit the fact that correlations in the connectivity matrix can modify the shape of the eigenvalue spectrum. Symmetrising (or anti-symmetrising) \mathbf{W} has been shown to generate elliptical (as opposed to circular) eigenspectra, in the case of “centered” matrices where the distinction between excitatory and inhibitory neurons is not made [13]. Here we consider a modification of the sparse neural matrices studied in section 3 that achieves this slimming effect in the case of balanced networks (see the insets in Fig. 6a). All non-zero entries assume a value $\pm w_0/\sqrt{N}$, the sign depending on the excitatory

⁶This happens provided the eigenvectors are sorted in decreasing order of their corresponding eigenvalue moduli, prior to going through the Gram-Schmidt orthonormalisation process. This results in a unique Schur basis.

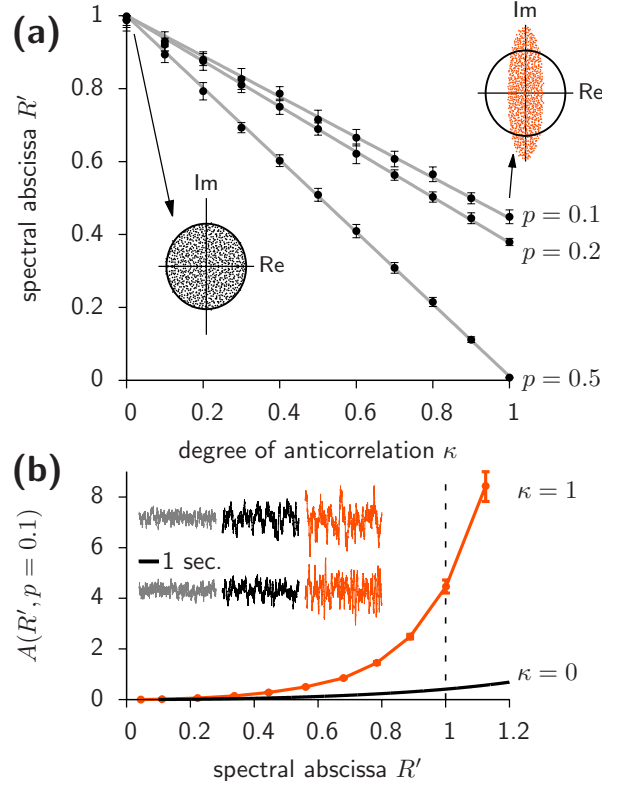


Figure 6: Example of a network structure that favors nonnormal amplification: unidirectional vs. bidirectional synaptic connections. **(a)** We varied the degree of anticorrelation between reciprocal weights in the connectivity matrix, as the fraction κ of the maximum value it can assume, which is dictated by the connection sparsity (see text). This caused the eigenspectrum to stretch more and more along the vertical axis (compare the two insets), effectively decreasing the spectral abscissa R' (black filled circles). Empirical data was obtained from numerically computing the eigenvalues of 20 different matrices of size $N = 500$. Errorbars denote standard deviations over all trials. Gray lines are linear fits. **(b)** Nonnormal amplification as a function of the spectral abscissa R' . When all connections between an excitatory (E) and an inhibitory (I) cell are made reciprocal, while all $E \rightarrow E$ and $I \rightarrow I$ connections are kept unidirectional (orange curve, corresponding to $\kappa = 1$ in (a)), stronger amplification is obtained in the fast dynamical regime ($R' \ll 1$). The black curve is here reproduced from Fig. 4 (purely random case, $\kappa = 0$) for comparison. The inset displays examples of 4-second snapshots of activity in a disconnected network (grey), a random network (black, $\kappa = 0$), and a maximally (though not fully) antisymmetric network (orange, $\kappa = 1$). The spectral abscissa was set to $R' = 0.9$. Traces were obtained from a direct simulation of Eq. 1, and are shown here only for two randomly chosen neurons.

versus inhibitory nature of the presynaptic neuron. Whether a connection exists (non-zero entry) is decided as follows. Connection w_{ij} with $i \geq j$ exists with probability p . If $i \neq j$, the reciprocal connection w_{ji} then exists with probability $p + c_{ij}(1 - p)$ if w_{ij} exists too, or with probability $p(1 - c_{ij})$ if it does not. In comparison to the random networks considered above (Eq. 3), this connectivity scheme preserves the mean weight $\bar{w} \equiv \langle w_{ij} \rangle = \pm p w_0 / \sqrt{N}$ as well as the weight

variance $\langle (w_{ij} - \bar{w})^2 \rangle = p(1-p)w_0^2/N$ while giving full control over their normalized covariance c_{ij} . Note that c_{ij} can assume positive values as high as $c_{\max} = 1$, in which case all connections are bidirectional. However, c cannot go below $c_{\min} = -p/(1-p)$, which stems from the sparsity condition that imposes a certain degree of symmetry in \mathbf{W} : because both w_{ij} and w_{ji} are zero with high probability, they will often be null together, meaning that they cannot be fully anti-correlated. The limit case $c = c_{\min}$ corresponds to the complete absence of reciprocal connections. Since we aim at tilting \mathbf{W} towards antisymmetry, we choose $c_{ij} = \kappa c_{\min}$ when neurons i and j are of the same type, and $c_{ij} = \kappa c_{\max}$ when the two neurons have different types. Thus $0 < \kappa < 1$ parameterises the degree of antisymmetry in \mathbf{W} . As can be seen in Fig. 6a, increasing κ effectively decreases the spectral abscissa $R' = \max_{\lambda} \text{Re}(\lambda)$, although it is designed not to affect the overall connectivity “power” $\sum w_{ij}^2$ which is the relevant quantity for amplification. Thus, for a fixed level of dynamical slowing (i.e. fixed R'), antisymmetric connectivity matrices can assume larger weight strengths and thereby yield stronger nonnormal amplification than their random counterparts, as depicted in Fig. 6b. Finally, note that a matrix with $\kappa = 1$ is *not* purely antisymmetric ($\mathbf{W}^\dagger \neq -\mathbf{W}$). In fact, neural connectivity matrices can never be fully antisymmetric, because of the constraint that neurons can be only excitatory or only inhibitory. This is an advantageous restriction here, because a fully antisymmetric matrix – just like a fully symmetric one – is in fact a normal operator that cannot support transient amplification.

8 Discussion

The nonnormal nature of the neuronal connectivity could play a major role in the functional dynamics of cortical networks. It can allow fast transients to develop along well-defined activity motifs stored in the pattern of synaptic efficacies. In networks with locally dense connectivity, the balance between excitation and inhibition has been shown to generate such type of amplification, accordingly termed “balanced amplification” [5]. We have revisited this feature in sparse balanced networks in which any two neurons are connected randomly with some probability. Random networks had already been studied in terms of their pseudospectrum [8], which only provides bounds on amplification. We have chosen a more direct approach and assessed nonnormality in terms of its functional impact in networks driven by stochastic external input. We have explicitly calculated the strength of the activity fluctuations that can only be attributed to the nonnormality of the recurrent connectivity. We found nonnormal amplification to be very weak, concluding that the only way to obtain large amplification in random networks is to allow for significant dynamical slowing. If the dynamics are to be kept fast, then the connectivity needs some structuring, so as to allow synaptic weights to take up larger values and to discourage the emergence of large positive eigenvalues. We have given an example of min-

imal network structure, namely connection antisymmetry, that achieves precisely this. More adaptive ways of shaping the connectivity, such as synaptic plasticity, could also be considered. In particular, inhibitory synaptic plasticity has recently been shown to suppress the attractor dynamics of a few activity motifs embedded in a spiking network, while still permitting their transient recall [12].

Nonnormal amplification could provide a mechanistic account for the often reported transient nature of both spontaneous and evoked activity in primary sensory cortices. Moreover, from a functional viewpoint, amplification without slowing could be a highly relevant feature in areas involved in the processing of fast-changing signals. If transient amplification by the synaptic connectivity is meant to allow past experience to be reflected in the responses to sensory stimuli (see e.g. [14]), then it is quite reassuring that random networks are poor amplifiers, for it implies that nothing can be amplified that has not been learned.

Here we have focused on spontaneous activity, i.e. on the fluctuations elicited by isotropic external noise that is totally uninformed of the frozen structure of the connectivity matrix. The equivalent triangular form of a nonnormal connectivity matrix suggests that neuronal networks should be more sensitive along some input directions than along others, so they could still respond vigorously though transiently to some carefully chosen input patterns (evoked activity). The first Schur mode, for example, is indeed such a preferred pattern [6]. This anisotropy prompts two important questions. First, how many different (orthogonal?) directions of high sensitivity does a network possess? Similarly, in how many distinguishable directions can the network amplify those preferred input signals? These quantities taken together could define the “nonnormal information capacity” of a network, reminiscent of the concept of memory capacity in attractor networks.

Finally, our analysis has revealed that the nonnormality of balanced networks is to a large extent reflected in large “DC” fluctuations. This seems to be a general feature of networks in which neurons can either be excitatory or inhibitory, but not of a mixed type [15]. It is somewhat disappointing that however strong activity fluctuations are in individual neurons, they always comprise a finite fraction of common variability. This is because the variance of the overall population activity is of the same order as the activity variance of the individual neurons (Eq. 16). Should computations exploit the fluctuations along the remaining $N - 1$ degrees of freedom of the network, complications in decoding the current network state would most certainly arise from a single dimension dominating the dynamics. However, we wish to point out that these large DC fluctuations are in fact a direct consequence of the *exact* excitation-inhibition balance considered here. We show in Appendix C that when inhibition *dominates* over excitation, the variance of the population activity becomes suddenly inversely proportional to the network size. Furthermore, the mean pairwise correlation coefficient in the network scales similarly, and thus

vanishes in large networks unless the E-I balance is exact. Note that this phenomenon is *not* mediated by a destruction of the strong feedforward link from the global balance disruption \mathbf{d} onto the DC mode \mathbf{v} , as described at the end of section 3. Increasing the overall amount of inhibition does preserve this strong link, but cancels its amplifying effect by imposing an equally strong negative feedback from the DC mode onto itself (see Appendix C). This dynamic cancellation of fluctuations and correlations was already shown to arise in balanced networks of spiking neurons [16]. Our results obtained for linear networks therefore suggest it may be a very general feature of inhibition-dominated balanced networks, and that fine-tuning the balance until it becomes exact [12] may strongly affect the dynamics of the network and the resulting correlation structure.

Acknowledgments

G.H. was partially supported by the European Union Framework 7 ICT Project no. 215910 (BIOTACT, www.biotact.org). T.P.V. was supported by the European Community's Seventh Framework Marie Curie International Reintegration grant no. 268436. G.H and T.P.V. were also partially supported by BrainScaleS, an EU FET-Proactive FP7 project under grant no. 269921. Thanks to L. Abbott, J. de la Rocha, K. Miller and Y. Ahmadian for helpful discussions.

A Amplification in random triangular networks

In this appendix we derive an exact expression for amplification in random strictly triangular networks with linear stochastic dynamics as in Eq. 1, where the non-zero elements of the coupling matrix \mathbf{T} are drawn from an arbitrary distribution with zero mean and variance α^2/N where N is the network size. Though no closed-form solution is known for the zero time lag covariance matrix Σ , we know from the theory of multidimensional Ornstein-Uhlenbeck processes that it satisfies the so-called Lyapunov equation [17]

$$(\mathbf{T} - \mathbb{1})\Sigma + \Sigma(\mathbf{T}^\dagger - \mathbb{1}) = -\tau\sigma_\xi^2\mathbb{1} \quad (20)$$

Equating component $(i, j < i)$ on both sides of Eq. 20 yields:

$$\sigma_{ij} = \frac{1}{2} \sum_{k=1}^{i-1} t_{ik}\sigma_{jk} + \frac{1}{2} \sum_{k=1}^{j-1} t_{jk}\sigma_{ik} \quad (21)$$

and equating the diagonal term (i, i) on both sides gives the variance of Schur mode i :

$$\sigma_{ii} = \frac{\tau\sigma_\xi^2}{2} + \sum_{j=1}^{i-1} t_{ij}\sigma_{ij} \quad (22)$$

Combining Eqs. 21 and 22 yields

$$\sigma_{ii} = \frac{\tau\sigma_\xi^2}{2} + \frac{1}{2} \sum_{j=1}^{i-1} t_{ij} \left(\sum_{k=1}^{i-1} t_{ik}\sigma_{jk} + \sum_{k=1}^{j-1} t_{jk}\sigma_{ik} \right) \quad (23)$$

in which σ_{jk} and σ_{ik} are to be recursively obtained from Eq. 21 with proper replacement of indices. We would like to calculate the expected value over the t_{ij} coefficients, i.e. over multiple realisations of random matrix \mathbf{T} . Explicitly expanding the sums will reveal cross-terms like $\langle t_{ij}t_{k\ell} \rangle$. Those vanish if $i \neq k$ or $j \neq \ell$, because the coupling coefficients are taken to be uncorrelated. The only remaining terms will be powers of the variance α^2/N . Here we seek a truncation to order α^4/N^2 . Let us calculate:

$$\begin{aligned} \langle \sigma_{ii} \rangle &= \frac{\tau\sigma_\xi^2}{2} + \frac{1}{2} \sum_{j=1}^{i-1} \sum_{k=1}^{i-1} \langle t_{ij}t_{jk}\sigma_{jk} \rangle \\ &\quad + \frac{1}{2} \sum_{j=1}^{i-1} \sum_{k=1}^{j-1} \langle t_{ij}t_{jk}\sigma_{ik} \rangle \end{aligned} \quad (24)$$

Because the network of Schur modes is purely feedforward, the cross-covariance σ_{jk} for $(j, k) < i$ is independent of the coupling coefficients t_{ij} and t_{ik} , thus $\langle t_{ij}t_{jk}\sigma_{jk} \rangle = \langle t_{ij}t_{jk} \rangle \langle \sigma_{jk} \rangle$. The only non-vanishing term in the first double-sum is therefore obtained for $k = j$, giving

$$\langle \sigma_{ii} \rangle = \frac{\tau\sigma_\xi^2}{2} + \frac{\alpha^2}{2N} \sum_{j=1}^{i-1} \langle \sigma_{jj} \rangle + \frac{1}{2} \sum_{j=1}^{i-1} \sum_{k=1}^{j-1} \langle t_{ij}t_{jk}\sigma_{ik} \rangle \quad (25)$$

Let us expand the expression in the second double-sum using Eq. 21:

$$\begin{aligned} \langle t_{ij}t_{jk}\sigma_{ik} \rangle &= \frac{1}{2} \sum_{\ell=1}^{i-1} \langle t_{ij}t_{jk}t_{i\ell}\sigma_{k\ell} \rangle \\ &\quad + \frac{1}{2} \sum_{\ell=1}^{k-1} \langle t_{ij}t_{jk}t_{k\ell}\sigma_{i\ell} \rangle \end{aligned} \quad (26)$$

As above, the first sum vanishes except for $\ell = j$. Should one continue and expand the second sum, one would receive terms of order α^6/N^3 and more which are discarded here (see above). Hence

$$\langle t_{ij}t_{jk}\sigma_{ik} \rangle = \frac{\alpha^2}{2N} \langle t_{jk}\sigma_{jk} \rangle + \dots \quad (27)$$

Using similar arguments, we expand $\langle t_{jk}\sigma_{jk} \rangle$ to order α^2/N and receive:

$$\langle t_{jk}\sigma_{jk} \rangle = \frac{\alpha^2}{2N} \langle \sigma_{kk} \rangle + \dots \quad (28)$$

From Eq. 25 it therefore follows that

$$\begin{aligned} \langle \sigma_{ii} \rangle &= \frac{\tau\sigma_\xi^2}{2} + \frac{\alpha^2}{2N} \sum_{j=1}^{i-1} \langle \sigma_{jj} \rangle \\ &\quad + \frac{\alpha^4}{8N^2} \sum_{j=1}^{i-1} \sum_{k=1}^{j-1} \langle \sigma_{kk} \rangle \end{aligned} \quad (29)$$

Defining $f_i = 2\langle \sigma_{ii} \rangle / (\sigma_\xi^2\tau)$, we end up with a recursive equation for the build-up of relative variance down the feedforward network of Schur modes:

$$f_i = 1 + \frac{\alpha^2}{2N} \sum_{j=1}^{i-1} f_j + \frac{\alpha^4}{8N^2} \sum_{j=1}^{i-1} \sum_{k=1}^{j-1} f_k \quad (30)$$

Now we define $x = i/N$ (thus $0 \leq x \leq 1$) and rewrite Eq. 30 as

$$f_{xN} = 1 + \frac{\alpha^2 x}{2i} \sum_{j=1}^{i-1} g\left(\frac{xj}{i}\right) + \frac{\alpha^4 x^2}{8i^2} \sum_{j=1}^{i-1} \sum_{k=1}^{j-1} g\left(\frac{xk}{i}\right) \quad (31)$$

In the limit $N \rightarrow \infty$ with constant $x = i/N$ ratio, the sums on the r.h.s. converge to their corresponding Riemann integrals, endowing f_{xN} with a proper limit $g(x)$:

$$g(x) = 1 + \frac{\alpha^2 x}{2} \int_0^1 g(xs) ds + \frac{\alpha^4 x^2}{8} \int_0^1 ds \int_0^1 ds' \Theta(s-s') g(xs') \quad (32)$$

where Θ is the Heaviside function. This convergence stems from the $1/N$ scaling of the variance α^2/N . Using straightforward changes of variables ($s \mapsto s/x$), we end up with an integral equation for g , the continuous variance profile along the (now infinitely large) network of Schur patterns:

$$g(x) = 1 + \frac{\alpha^2}{2} \int_0^x g(s) ds + \frac{\alpha^4}{8} \int_0^x ds \int_0^s ds' g(s') \quad (33)$$

Differentiating Eq. 33 twice with respect to x yields a second-order differential equation for g

$$g''(x) = \frac{\alpha^2}{2} g'(x) + \frac{\alpha^4}{8} g(x) \quad (34)$$

with initial conditions $g(0) = 1$, $g'(0) = \alpha^2/2$, and $g''(0) = 3\alpha^4/8$. The solution is precisely $g^{\text{LB}}(x)$ given in Eq. 12 of the main text. It is only a lower-bound on the true variance profile $g(x)$ since all the higher-order terms in α^2 that we have neglected are positive. This approximation proves reasonable for $\alpha^2 < 3$ as shown in Fig. 3a (dashed blue lines). Further integrating over x yields a lower-bound on nonnormal amplification $A_0(\alpha^2) \equiv \int_0^1 g(x) dx - 1$ (Fig. 3b, dashed blue line):

$$A_0^{\text{LB}}(\alpha^2) = \frac{2}{\alpha^2 \sqrt{3}} \exp\left(-\frac{(\sqrt{3}-1)\alpha^2}{4}\right) \times \left[\exp\left(\frac{\sqrt{3}\alpha^2}{2}\right) - 1 \right] - 1 \quad (35)$$

Instead of truncating $\langle \sigma_{ii} \rangle$ to order α^4 , one can also decide to start again from Eq. 24 and keep all terms up to order n . This requires careful counting, and results in a differential equation of order n , reading

$$g^{(n)}(x) = \frac{\alpha^2}{2} \sum_{k=0}^n C_k \left(\frac{\alpha^2}{4}\right)^k g^{(n-k-1)}(x) \quad (36)$$

where $C_k = (2k)!/[k!(k+1)!]$ is the k^{th} Catalan number. Assuming $g(x)$ can be written for $0 \leq x \leq 1$ as a convergent power series

$$g(x) = \lim_{K \rightarrow \infty} \sum_{k=0}^K \beta_k x^k \quad (37)$$

and equating $g^{(k)}(0)$ in both Eqs. 36 and 37 yields the results of Eqs. 13 – 15.

B Variance of the DC component

The last Schur mode is fed by the activities of all previous Schur vectors, weighted by couplings with variance ζ_0^2/N . The same calculation that led to Eq. 30 in this case leads to

$$f_N = 1 + \frac{\zeta_0^2}{2} \sum_{j=1}^{N-1} f_j + \frac{\zeta_0^2 R^2}{8N} \sum_{j=1}^{N-1} \sum_{k=1}^{j-1} f_k + \dots \quad (38)$$

which can be rewritten as

$$\frac{f_N}{N} = \frac{1}{N} + \frac{\zeta_0^2}{R^2} \left(\frac{R^2}{2N} \sum_{j=1}^{N-1} f_j + \frac{R^4}{8N^2} \sum_{j=1}^{N-1} \sum_{k=1}^{j-1} f_k + \dots \right) \quad (39)$$

where the sums were previously calculated in the limit $N \rightarrow \infty$ (Eqs. 30 – 37). We thus recover

$$\lim_{N \rightarrow \infty} \frac{f_N}{N} = \frac{\zeta_0^2}{R^2} [g(1) - 1] \quad (40)$$

With ζ_0^2 given by Eq. 8 we arrive at Eq. 16 of the main text.

C Exactly balanced vs. inhibition-dominated networks

In this paper, we have considered connectivities in which weights were either zero or $\pm w_0/\sqrt{N}$, the \pm sign depending on the excitatory vs. inhibitory nature of the presynaptic neuron (Eq. 3). Furthermore, the number of cells of both types was identical. The total inhibitory synaptic strength thus exactly matched its excitatory counterpart. In this appendix, we wish to show that if the non-zero inhibitory weights are stronger, i.e. $-\gamma w_0/\sqrt{N}$ with $\gamma > 1$, the dynamics of the overall population activity is strongly affected.

We have seen that the “DC” mode $\mathbf{v} = (1, \dots, 1)/\sqrt{N}$ is an eigenvector of \mathbf{W} . Let λ_v denote the associated eigenvalue, which quantifies the effective decay rate of the DC component in the network of neurons. If the E-I balance is exact ($\gamma = 1$) as assumed throughout the paper, then $\lambda_v = 0$. More generally, however, one can calculate

$$\lambda_v = -\frac{p w_0 (\gamma - 1)}{2} \cdot \sqrt{N} \quad (41)$$

We see there is an unexpected scaling that the exact balance was hiding: $-\lambda_v \sim \mathcal{O}(\sqrt{N})$. Note that all other eigenvalues are now scattered inside the disk of radius

$$R = w_0 \sqrt{\frac{(1 + \gamma^2)p(1 - p)}{2}} \quad (42)$$

though no longer uniformly so since the variance of the inhibitory and excitatory weights now differ by a factor of γ^2 [10]. Having kept the focus of this paper on nonnormal effects, we have intentionally set aside the contributions of the eigenvalues to the overall amplification in the network. When $\lambda_v = 0$ (perfect balance), our prediction that the average population activity $\mu(t) \equiv \sum x_i(t)/N$ should have a variance of order $\mathcal{O}(1)$ was justified: the last Schur unit corresponding to this DC indeed receives $N - 1$ contributions of order $\mathcal{O}(1)$, and its decay time constant is simply $\tau \sim \mathcal{O}(1)$, yielding $\text{var}[\mu(t)] \sim \mathcal{O}(1)$. When inhibition dominates ($\gamma > 1$), the DC component suppresses itself via a negative feedback that scales with \sqrt{N} , yielding a very short decay time constant $\tau/(1 - \lambda_v) \sim \mathcal{O}(1/\sqrt{N})$ whose deviation from τ can no longer be neglected. To see what the implications of this scaling are for the variance of $\mu(t)$, let us reduce the dynamics of the DC to the following set of N stochastic differential equations:

$$\begin{aligned} dy_i &= -\frac{dt}{\tau} y_i + \sqrt{\frac{2}{\tau}} d\xi_i \quad \text{for } 1 \leq i < N \\ dy_N &= \frac{dt}{\tau} \left(-(1 - \lambda_v) y_N + \sum_{i=1}^{N-1} \varepsilon_i x_i \right) + \sqrt{\frac{2}{\tau}} d\xi_N \end{aligned} \quad (43)$$

Here y_1, \dots, y_{N-1} model the first $N - 1$ Schur units independently, with the appropriate noise terms such that they achieve a variance of one (corresponding to the limit of small amplification). They feed y_N – which models the activity of the last Schur unit, i.e. the DC component $\mu(t)\sqrt{N}$ – with couplings ε_i such that $\sum \varepsilon_i^2/N = \zeta_0^2$. We calculate the coupling variance ζ_0^2 the same way we did in section 3:

$$\zeta_0^2 = \frac{p^2 w_0^2 (1 + \gamma^2)}{2} \quad (44)$$

The variance $\text{var}[\mu(t)]$ of the overall neuronal population activity, here modeled by $\mu(t) \approx y_N(t)/\sqrt{N}$, is given by standard Ornstein-Uhlenbeck theory:

$$\text{var}(\mu(t)) = \frac{1}{N(1 - \lambda_v)} \left[1 + \frac{N\zeta_0^2}{2 - \lambda_v} \right] \quad (45)$$

Although we have neglected amplification and correlations in the first $N - 1$ Schur units, Eq. 45 does provides a good intuition for how the mean population activity $\mu(t) = y_N(t)/\sqrt{N}$ scales with the network size N , and provides a good qualitative match to numerical results even for a non-negligible spectral radius $R = 0.5$ (Fig. 7).

The asymptotics of $\text{var}[\mu(t)]$ are given by

$$\text{var}[\mu(t)] \sim \begin{cases} \frac{p^2 w_0^2}{2} & \text{if } \gamma = 1 \\ \frac{2(1 + \gamma^2)}{N(\gamma - 1)^2} & \text{if } \gamma > 1 \end{cases} \quad (46)$$

Thus, when inhibition dominates over excitation ($\gamma > 1$), the fluctuations of the overall population activity vanish for large networks, which was already shown in [16] for




Figure 7: Temporal fluctuations of the overall population firing rate in a balanced neuronal network. The variance of the average population activity $\mu(t) = \sum x_i(t)/N$ is reported as a function of the network size N , in logarithmic scale. When inhibition perfectly balances excitation ($\gamma = 1$), the variance is asymptotically independent of the network size (gray). When inhibition dominates ($\gamma > 1$), it scales with $1/N$ (black). The solid lines denote the approximation in Eq. 45. The dashed lines indicate the asymptotics (Eq. 46). Points denote the empirical variance obtained by simulating Eq. 1 for 100 seconds, for a neuronal network constructed as specified in section 3 with connectivity density $p = 0.1$. The spectral radius was set to $R = 0.1$ (top plot) and $R = 0.5$ (bottom plot).

inhibition-dominated networks of spiking neurons. In contrast, fine tuning the connectivity such that the balance becomes exact ($\gamma = 1$) opens the possibility for these fluctuations to subsist in arbitrarily large networks. This has profound consequences for the mean pairwise correlation $\bar{r} \equiv \sum_{i \neq j} \text{cov}[x_i(t), x_j(t)]/N^2$, as seen from the following identity

$$\bar{r} = \text{var}[\mu(t)] - \frac{1}{N^2} \sum_i \text{var}[x_i(t)] \quad (47)$$

We have seen that the average variance $\text{var}[x_i(t)]$ in the individual neurons (i.e. amplification as we define it) is $\mathcal{O}(1)$. Thus, Eq. 47 implies that \bar{r} scales with N in the same way $\text{var}[\mu(t)]$ does: either $\mathcal{O}(1)$ if the balance is perfect, or $\mathcal{O}(1/N)$ if inhibition dominates.

References

- [1] M. Tsodyks, T. Kenet, A. Grinvald, and A. Arieli. Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science*, 286:1943–1946, 1999.
- [2] T. Kenet, D. Bibitchkov, M. Tsodyks, A. Grinvald, and A. Arieli. Spontaneously emerging cortical representations of visual attributes. *Nature*, 425:954–956, 2003.
- [3] A. Luczak, P. Barthó, and K. D. Harris. Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*, 62:413–425, 2009.
- [4] J. A Goldberg, U. Rokni, and H. Sompolinsky. Patterns of ongoing activity and the functional architecture of the primary visual cortex. *Neuron*, 42(3):489–500, 2004.
- [5] Brendan K. Murphy and Kenneth D. Miller. Balanced amplification: A new mechanism of selective amplification of neural activity patterns. *Neuron*, 61(4):635–648, 2009.
- [6] S. Ganguli, D. Huh, and H. Sompolinsky. Memory traces in dynamical systems. *PNAS*, 105(48):18970–18975, 2008.
- [7] Mark S. Goldman. Memory without feedback in a neural network. *Neuron*, 61(4):621–634, 2009.
- [8] L N Trefethen and M Embree. *Spectra and pseudospectra: the behavior of nonnormal matrices and operators*. Princeton University Press, 2005.
- [9] Surya Ganguli and Peter Latham. Feedforward to the past: The relation between neuronal connectivity, amplification, and short-term memory. *Neuron*, 61(4):499–501, 2009.
- [10] K. Rajan and L. F. Abbott. Eigenvalue spectra of random matrices for neural networks. *Phys. Rev. Lett.*, 97:188104, 2006.
- [11] T. Tao. Outliers in the spectrum of iid matrices with bounded rank perturbations. *Probability Theory and Related Fields*, Online First:1–33, 2011. 10.1007/s00440-011-0397-9.
- [12] T. P. Vogels, H. Sprekeler, F. Zenke, C. Clopath, and W. Gerstner. Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science*, 334:1569, 2011.
- [13] H. J. Sommers, A. Crisanti, H. Sompolinsky, and Y. Stein. Spectrum of large random asymmetric matrices. *Phys. Rev. Lett.*, 60:1895–1898, 1988.
- [14] J. Fiser, P. Berkes, G. Orbán, and M. Lengyel. Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*, 14:119–130, 2010.
- [15] B. Kriener, T. Tetzlaff, A. Aertsen, M. Diesmann, and S. Rotter. Correlations and population dynamics in cortical networks. *Neural Comput.*, 20:2185–2226, 2008.
- [16] A. Renart, J. de la Rocha, P. Bartho, L. Hollender, N. Parga, A. Reyes, and K.D. Harris. The asynchronous state in cortical circuits. *Science*, 327:587, 2010.
- [17] C. W. Gardiner. *Handbook of stochastic methods: for physics, chemistry, and the natural sciences*. Berlin: Springer, 1985.